

Self-Shuffling Words

Émilie Charlier¹, Teturo Kamae², Svetlana Puzynina^{3,5}, and Luca Q. Zamboni^{4,5}

¹ Département de Mathématique, Université de Liège, Belgium
echarlier@ulg.ac.be

² Advanced Mathematical Institute, Osaka City University, Japan
kamae@apost.plala.or.jp

³ Sobolev Institute of Mathematics, Novosibirsk Russia
svepuz@utu.fi

⁴ Institut Camille Jordan, Université Lyon 1, France
lupastis@gmail.com

⁵ FUNDIM, University of Turku, Finland

Abstract. In this paper we introduce and study a new property of infinite words which is invariant under the action of a morphism: We say an infinite word $x \in \mathbb{A}^{\mathbb{N}}$, defined over a finite alphabet \mathbb{A} , is self-shuffling if x admits factorizations: $x = \prod_{i=1}^{\infty} U_i V_i = \prod_{i=1}^{\infty} U_i = \prod_{i=1}^{\infty} V_i$ with $U_i, V_i \in \mathbb{A}^+$. In other words, there exists a shuffle of x with itself which reproduces x . The morphic image of any self-shuffling word is again self-shuffling. We prove that many important and well studied words are self-shuffling: This includes the Thue-Morse word and all Sturmian words (except those of the form aC where $a \in \{0, 1\}$ and C is a characteristic Sturmian word). We further establish a number of necessary conditions for a word to be self-shuffling, and show that certain other important words (including the paper-folding word and infinite Lyndon words) are not self-shuffling. In addition to its morphic invariance, which can be used to show that one word is not the morphic image of another, this new notion has other unexpected applications: For instance, as a consequence of our characterization of self-shuffling Sturmian words, we recover a number theoretic result, originally due to Yasutomi, which characterizes pure morphic Sturmian words in the orbit of the characteristic.

1 Introduction

Let \mathbb{A} be a finite non-empty set. We denote by \mathbb{A}^* the set of all finite words $u = x_1 x_2 \dots x_n$ with $x_i \in \mathbb{A}$. The quantity n is called the length of u and is denoted $|u|$. The empty word, denoted ε , is the unique element in \mathbb{A}^* with $|\varepsilon| = 0$. We set $\mathbb{A}^+ = \mathbb{A} - \{\varepsilon\}$. We denote by $\mathbb{A}^{\mathbb{N}}$ the set of all one-sided infinite words $x = x_0 x_1 x_2 \dots$ with $x_i \in \mathbb{A}$.

Given k -finite or infinite words $x^{(1)}, x^{(2)}, \dots, x^{(k)} \in \mathbb{A}^* \cup \mathbb{A}^{\mathbb{N}}$ we denote by

$$\mathcal{S}(x^{(1)}, x^{(2)}, \dots, x^{(k)}) \subset \mathbb{A}^* \cup \mathbb{A}^{\mathbb{N}}$$

the collection of all words z for which there exists a factorization

$$z = \prod_{i=0}^{\infty} U_i^{(1)} U_i^{(2)} \dots U_i^{(k)}$$

with each $U_i^{(j)} \in \mathbb{A}^*$ and with $x^{(j)} = \prod_{i=0}^{\infty} U_i^{(j)}$ for $1 \leq j \leq k$. Intuitively, z may be obtained as a *shuffle* of the words $x^{(1)}, x^{(2)}, \dots, x^{(k)}$. In case $x^{(1)}, x^{(2)}, \dots, x^{(k)} \in \mathbb{A}^*$, each of the above products can be taken to be finite.

Finite word shuffles were extensively studied in [5]. Given $x \in \mathbb{A}^*$, it is generally a difficult problem to determine whether there exists $y \in \mathbb{A}^*$ such that $x \in \mathcal{S}(y, y)$ (see Open Problem 4 in [5]). However, in the context of infinite words, this question is essentially trivial: In fact, it is readily verified that if $x \in \mathbb{A}^{\mathbb{N}}$ is such that each $a \in \mathbb{A}$ occurring in x occurs an infinite number of times in x , then there exist infinitely many $y \in \mathbb{A}^{\mathbb{N}}$ with $x \in \mathcal{S}(y, y)$. Instead, in the framework of infinite words, a far more delicate question is the following:

Question 1. Given $x \in \mathbb{A}^{\mathbb{N}}$, does there exist an integer $k \geq 2$ such that $x \in \mathcal{S}(\underbrace{x, x, \dots, x}_k)$?

If such a k exists, we say x is *k-self-shuffling*.

Given $x = x_0 x_1 x_2 \dots \in \mathbb{A}^{\mathbb{N}}$ and an infinite subset $N = \{N_0 < N_1 < N_2 < \dots\} \subseteq \mathbb{N}$, we put $x[N] = x_{N_0} x_{N_1} x_{N_2} \dots \in \mathbb{A}^{\mathbb{N}}$. Alternatively,

Definition 1. For $x \in \mathbb{A}^{\mathbb{N}}$ and $k = 2, 3, \dots$, we say x is *k-self-shuffling* if there exists a k -element partition $\mathbb{N} = \bigcup_{i=1}^k N^i$ with $x[N^i] = x$ for each $i = 1, \dots, k$.

In case $k = 2$, we say simply x is *self-shuffling*. We note that if x is k -self-shuffling, then x is ℓ -self-shuffling for each $\ell \geq k$ but not conversely (see §2), whence each self-shuffling word is k -self-shuffling for all $k \geq 2$. In this paper we are primarily interested in self-shuffling words, however, many of the results presented here extend to general k . Thus $x \in \mathbb{A}^{\mathbb{N}}$ is self-shuffling if and only if x admits factorizations

$$x = \prod_{i=1}^{\infty} U_i V_i = \prod_{i=1}^{\infty} U_i = \prod_{i=1}^{\infty} V_i$$

with $U_i, V_i \in \mathbb{A}^+$.

The property of being self-shuffling is an intrinsic property of the word (and not of the associated language) and seems largely independent of its complexity (examples exist from the lowest to the highest possible complexity). The simplest class of self-shuffling words consists of all (purely) periodic words $x = u^{\omega}$. It is clear that if x is self-shuffling, then every letter $a \in \mathbb{A}$ occurring in x must occur an infinite number of times. Thus for instance, the ultimately periodic word 01^{ω} is not self-shuffling. As we shall see, many well-known words which are of interest in both combinatorics on words and symbolic dynamics, are self-shuffling. This includes for instance the famous *Thue-Morse* word

$$\mathbf{T} = 0110100110010110100101100110100110010110\dots$$

whose origins go back to the beginning of the last century with the works of the Norwegian mathematician Axel Thue [12,13]. The n th entry t_n of \mathbf{T} is defined as the sum modulo 2 of the digits in the binary expansion of n . While the Thue-Morse word appears naturally in many different areas of mathematics (from discrete mathematics to number theory to differential geometry-see [1] or [2]), proving that Thue-Morse is self-shuffling is somewhat more involved than expected.

Sturmian words constitute another important class of aperiodic self-shuffling words. Sturmian words are infinite words over a binary alphabet having exactly $n+1$ factors of length n for each $n \geq 0$ [8]. Their origin can be traced back to the astronomer J. Bernoulli III in 1772. They arise naturally in many different areas of mathematics including combinatorics, algebra, number theory, ergodic theory, dynamical systems and differential equations. Sturmian words are also of great importance in theoretical physics and in theoretical computer science and are used in computer graphics as digital approximation of straight lines. We show that all Sturmian words are self-shuffling except those of the form aC where $a \in \{0,1\}$ and C is a characteristic Sturmian word. Thus for every irrational number α , all (uncountably many) Sturmian words of slope α are self-shuffling except for two. Our proof relies on a geometric characterization of Sturmian words via irrational rotations on the circle.

So while there are many natural examples of aperiodic self-shuffling words, the property of being self-shuffling is nevertheless quite restrictive. We obtain a number of necessary (and in some cases sufficient) conditions for a word to be self-shuffling. For instance, if a word x is self-shuffling, then x begins in only finitely many Abelian border-free words. As an application of this we show that the well-known *paper folding word* is not self-shuffling. Infinite Lyndon words (i.e., infinite words which are lexicographically smaller than each of its suffixes) are also shown not to be self-shuffling.

One important feature of self-shuffling words stems from its invariance under the action of a morphism: The morphic image of a self-shuffling word is again self-shuffling. In some instances this provides a useful tool for showing that one word is not the morphic image of another. So for instance, the paper folding word is not the morphic image of any self-shuffling word. However this application requires knowing a priori whether a given word is or is not self-shuffling. In general, to show that a word is self-shuffling, one must actually exhibit a shuffle. Self-shuffling words have other unexpected applications particularly in the study of fixed points of substitutions. For instance, as an almost immediate consequence of our characterization of self-shuffling Sturmian words, we recover a result, first proved by Yasutomi via number theoretic methods, which characterizes pure morphic Sturmian words in the orbit of the characteristic.

2 Examples & Non-examples

In this section we list some examples and non-examples of self-shuffling words. As usual in combinatorics on words, we follow notation from [8].

Fibonacci word: The Fibonacci infinite word

$$x = 0100101001001010010100\dots$$

is defined as the fixed point of the morphism φ given by $0 \mapsto 01, 1 \mapsto 0$. It is readily verified that $\varphi^2(a) = \varphi(a)a$ for each $a \in \{0, 1\}$. Whence, writing $x = x_0x_1x_2\dots$ with each $x_i \in \{0, 1\}$ we obtain

$$\begin{aligned} x &= x_0x_1x_2\dots = \varphi(x_0)\varphi(x_1)\varphi(x_2)\dots = \varphi^2(x_0)\varphi^2(x_1)\varphi^2(x_2)\dots = \\ &= \varphi(x_0)x_0\varphi(x_1)x_1\varphi(x_2)x_2\dots \end{aligned}$$

which shows that x is self-shuffling. In contrast, the word $y = 0x$ is not self-shuffling. The word y starts with infinitely many prefixes of the form $0B1$ with B a palindrome. It follows that $0B1$ is *Abelian border-free* (i.e., no proper suffix of $0B1$ is Abelian equivalent to a proper prefix of U). By Proposition 3 the word y is not self-shuffling.

Paper-folding word: The paper-folding word

$$x = 00100110001101100010\dots$$

is a Toeplitz word generated by the pattern $u = 0?1?$ (see[4]). It is readily verified that x begins in arbitrarily long Abelian border-free words and hence by Proposition 3 is not self-shuffling. More precisely, the prefixes u_j of x of length $n_j = 2^j - 1$ are Abelian border-free. Indeed, it is verified that for each $k < n_j$, we have $|\text{pref}_k(u_j)|_0 > k/2$ while $|\text{suff}_k(u_j)|_0 \leq k/2$. Here $\text{pref}_k(u)$ (resp., $\text{suff}_k(u)$) denotes the prefix (resp., suffix) of length k of a word u , an $|u|_a$ is the number of occurrences of a letter a in u .

A 3-self-shuffling word which is not self-shuffling: Let y denote the fixed point of the morphism $\sigma : 0 \mapsto 0001$ and $1 \mapsto 0101$, and put

$$x = 0^{-2}y = 0100010001010100010001000101010001000101010001010100\dots$$

Then for each prefix u_j of x of length $4^j - 2$, the longest Abelian border of u_j of length less than or equal to $(4^j - 2)/2$ has length 2. Hence x is not self-shuffling (see Proposition 3). The 3-shuffle is given by the following:

$$\begin{aligned} U_0 &= 0100, U_1 = 01, & \dots, U_{4i+2} &= \varepsilon, & U_{4i+3} &= \sigma^{i+1}(0100), \\ & & U_{4i+4} &= \sigma(0), & U_{4i+5} &= (\sigma(0))^{-1}\sigma^{i+1}(01), \\ V_0 &= 0100, V_1 = 01, & \dots, V_{4i+2} &= (\sigma(0))^{-1}\sigma^{i+1}(0), & V_{4i+3} &= \sigma(0), \\ & & V_{4i+4} &= (\sigma(0))^{-1}\sigma^{i+1}(01)\sigma(0), & V_{4i+5} &= \varepsilon, \\ W_0 &= 01, W_1 = (\sigma(0))^2, \dots, & W_{4i+2} &= \varepsilon, & W_{4i+3} &= (\sigma(0))^{-1}\sigma^{i+1}(01), \\ & & W_{4i+4} &= \varepsilon, & W_{4i+5} &= \sigma^{i+2}(0)\sigma(0). \end{aligned}$$

It is then verified that

$$x = \prod_{i=0}^{\infty} U_i V_i W_i = \prod_{i=0}^{\infty} U_i = \prod_{i=0}^{\infty} V_i = \prod_{i=0}^{\infty} W_i,$$

from which it follows that x is 3-self-shuffling.

A recurrent binary self-shuffling word with full complexity: For each positive integer n , let z_n denote the concatenation of all words of length n in increasing lexicographic order. For example, $z_2 = 00011011$. For $i \geq 0$ put

$$v_i = \begin{cases} z_n, & \text{if } i = n2^{n-1} \text{ for some } n, \\ 0^i 1^i, & \text{otherwise,} \end{cases}$$

and define

$$x = \prod_{i=0}^{\infty} X_i = 010100110^3 011^3 0^4 010^2 1^2 011^4 \dots,$$

where $X_0 = X_1 = 01$, $X_2 = 0011$, and for $i \geq 3$, $X_i = 0^i y_{i-2} 1^i$, where $y_{i-2} = y_{i-3} v_{i-2} y_{i-3}$, and $y_0 = \varepsilon$. We note that x is recurrent and has full complexity (since it contains z_n as a factor for every n).

To show that the word x is self-shuffling, we first show that $X_{i+1} \in \mathcal{S}(X_i, X_i)$. Take $N = \{0, \dots, i-1, i+1, \dots, 2^i - i, 2^i - i + v_{i-1}|_1, 2^{i+1} - i - 1\}$, where $u|_1$ denotes the set of positions j of a word u in which the j -th letter u_j of u is equal to 1. Then it is straightforward to see that $X_i = X_{i+1}[N] = X_{i+1}[\{1, \dots, 2^{i+1}\} \setminus N]$. The self-shuffle of x is built in a natural way concatenating shuffles of X_i starting with $U_0 = V_0 = 01$, so that $X_0 \dots X_{i+1} \in \mathcal{S}(X_0 \dots X_i, X_0 \dots X_i)$.

3 General Properties

In this section we develop several fundamental properties of self-shuffling words. The next two propositions show the invariance of self-shuffling words with respect to the action of a morphism:

Proposition 1. *Let \mathbb{A} and \mathbb{B} be finite non-empty sets and $\tau : \mathbb{A} \rightarrow \mathbb{B}^*$ a morphism. If $x \in \mathbb{A}^{\mathbb{N}}$ is self-shuffling, then so is $\tau(x) \in \mathbb{B}^{\mathbb{N}}$.*

Proof. If $x \in \mathcal{S}(x, x)$, then we can write $x = \prod_{i=1}^{\infty} U_i V_i = \prod_{i=1}^{\infty} U_i = \prod_{i=1}^{\infty} V_i$. Whence $\tau(x) = \prod_{i=1}^{\infty} \tau(U_i V_i) = \prod_{i=1}^{\infty} \tau(U_i) \tau(V_i) = \prod_{i=1}^{\infty} \tau(U_i) = \prod_{i=1}^{\infty} \tau(V_i)$ as required.

Proposition 2. *Let $\tau : \mathbb{A} \rightarrow \mathbb{A}^*$ be a morphism, and $x \in \mathbb{A}^{\mathbb{N}}$ be a fixed point of τ .*

1. *Let u be a prefix of x and k be a positive integer such that $\tau^k(a)$ begins in u for each $a \in \mathbb{A}$. Then if x is self-shuffling, then so is $u^{-1}x$.*
2. *Let $u \in \mathbb{A}^*$, and let k be a positive integer such that $\tau^k(a)$ ends in u for each $a \in \mathbb{A}$. Then if x is self-shuffling, then so is ux .*

Proof. We prove only item(1) since the proof of (2) is essentially identical. Suppose $x = \prod_{i=1}^{\infty} U_i V_i = \prod_{i=1}^{\infty} U_i = \prod_{i=1}^{\infty} V_i$. Then by assumption, for each $i \geq 1$

we can write $\tau^k(U_i) = uU'_i$ and $\tau^k(V_i) = uV'_i$ for some $U'_i, V'_i \in \mathbb{A}^*$. Put $X_i = U'_i u$ and $Y_i = V'_i u$. Then since

$$x = \tau^k(x) = \prod_{i=1}^{\infty} \tau^k(U_i V_i) = \prod_{i=1}^{\infty} \tau^k(U_i) \tau^k(V_i) = \prod_{i=1}^{\infty} \tau^k(U_i) = \prod_{i=1}^{\infty} \tau^k(V_i),$$

we deduce that

$$u^{-1}x = \prod_{i=1}^{\infty} X_i Y_i = \prod_{i=1}^{\infty} X_i = \prod_{i=1}^{\infty} Y_i.$$

Corollary 1. *Let $\tau : \mathbb{A} \rightarrow \mathbb{A}^*$ be a primitive morphism, and $a \in \mathbb{A}$. Suppose $\tau(b)$ begins in a (respectively ends in a) for each letter $b \in \mathbb{A}$. Suppose further that the fixed point $\tau^\infty(a)$ is self-shuffling. Then every forward shift (respectively backward shift) of $\tau^\infty(a)$ is self-shuffling.*

Remark 1. Since the Fibonacci word is self-shuffling and is fixed by the primitive morphism $0 \mapsto 01, 1 \mapsto 0$, it follows from Corollary 1 that every tail of the Fibonacci word is self-shuffling.

There are a number of necessary conditions that a self-shuffling word must satisfy, which may be used to deduce that a given word is not self shuffling. For instance:

Proposition 3. *If $x \in \mathbb{A}^{\mathbb{N}}$ is self-shuffling, then for each positive integer N there exists a positive integer M such that every prefix u of x with $|u| \geq M$ has an Abelian border v with $|u|/2 \geq |v| \geq N$. In particular, x must begin in only a finite number of Abelian border-free words.*

Proof. Suppose to the contrary that there exist factorizations $x = \prod_{i=0}^{\infty} U_i V_i = \prod_{i=0}^{\infty} U_i = \prod_{i=0}^{\infty} V_i$ with $U_i, V_i \in \mathbb{A}^+$, and there exists N such that for every M there exists a prefix u of x with $|u| \geq M$ which has no Abelian borders of length between N and $|u|/2$. Take $M = |\prod_{i=0}^N U_i V_i|$ and a prefix u satisfying these conditions. Then there exist non-empty proper prefixes U' and V' of u such that $u \in \mathcal{S}(U', V')$ with $|U'|, |V'| > N$. Writing $u = U'U''$ it follows that U'' and V' are Abelian equivalent. This contradicts that u has no Abelian borders of length between N and $|u|/2$.

An extension of this argument gives both a necessary and sufficient condition for self-shuffling in terms of Abelian borders (which is however difficult to check in practice). For $u \in \mathbb{A}^*$ let $\Psi(u)$ denote the *Parikh vector* of u , i. e., $\Psi(u) = (|u|_a)_{a \in \mathbb{A}}$, where $|u|_a$ denotes the number of occurrences of a in u .

Definition 2. Given $x \in \mathbb{A}^{\mathbb{N}}$, we define a directed graph $G_x = (V_x, E_x)$ with vertex set

$$V_x = \{(n, m) \in \mathbb{N}^2 \mid \Psi(\text{pref}_n x) + \Psi(\text{pref}_m x) = \Psi(\text{pref}_{n+m} x)\}$$

and the edge set

$$E_x = \{ ((n, m), (n', m')) \in V_x \times V_x \mid \\ n' = n + 1 \text{ and } m' = m \text{ or } m' = m + 1 \text{ and } n' = n \}.$$

We say that G_x connects $\mathbf{0}$ to ∞ if there exists an infinite path $\prod_{j=1}^{\infty} (n_j, m_j)$ in G_x such that $(n_0, m_0) = (0, 0)$ and $n_j, m_j \rightarrow \infty$ as $j \rightarrow \infty$.

Theorem 1. *A word $x \in \mathbb{A}^{\mathbb{N}}$ is self-shuffling if and only if the graph G_x connects $\mathbf{0}$ to ∞ .*

The theorem gives a constructive necessary and sufficient condition for self-shuffling since a path to infinity defines a self-shuffle.

As we shall now see, lexicographically extremal words are never self-shuffling. Let (\mathbb{A}, \leq) be a finite linearly ordered set. Then \leq induces a partial ordering \leq_{lex} on \mathbb{A}^+ and $\mathbb{A}^{\mathbb{N}}$ defined as follows: If $u, v \in \mathbb{A}^+$ (or $\mathbb{A}^{\mathbb{N}}$) we write $u \leq_{\text{lex}} v$ if either $u = v$ or if u is lexicographically smaller than v . In the latter case we write $u <_{\text{lex}} v$. If u is a proper prefix of v then u and v are not comparable.

Let $x \in \mathbb{A}^{\mathbb{N}}$. A factor u of x is called *minimal* (in x) if $u \leq_{\text{lex}} v$ for all factors v of x with $|v| = |u|$. An infinite word y in the shift orbit closure S_x of x is called *Lyndon* (in S_x) if every prefix of y is minimal in x . The proof of the following result is omitted for space considerations:

Theorem 2. *Let (A, \leq) be a linearly ordered finite set and let $x \in A^{\mathbb{N}}$. Let $y, z \in S_x$ with y Lyndon and aperiodic. Then for each $w \in \mathcal{S}(y, z)$, we have $w <_{\text{lex}} z$. In particular, taking $z = y$ we deduce that y is not self-shuffling.*

Let \mathbb{A} be a finite non-empty set. We say $x \in \mathbb{A}^{\mathbb{N}}$ is *extremal* if there exists a linear ordering \leq on \mathbb{A} with respect to which x is Lyndon. As an immediate consequence of Theorem 2 we obtain:

Corollary 2. *Let \mathbb{A} be a finite non-empty set and $x \in \mathbb{A}^{\mathbb{N}}$ be an aperiodic extremal word. Then x is not self-shuffling.*

Remark 2. Let $x = 11010011001011010010110\dots$ denote the first shift of the Thue-Morse infinite word. It is well known that x is extremal and hence is not self-shuffling; yet it can be verified that x begins in only a finite number of Abelian border-free words.

4 The Thue-Morse Word Is Self-Shuffling

Theorem 3. *The Thue-Morse word $\mathbf{T} = 011010011001\dots$ fixed by the morphism τ mapping $0 \mapsto 01$ and $1 \mapsto 10$ is self-shuffling.*

Proof. For $u \in \{0, 1\}^*$ we denote by \bar{u} the word obtained from u by exchanging 0s and 1s. Let $\sigma : \{1, 2, 3, 4\} \rightarrow \{1, 2, 3, 4\}^*$ be the morphism defined by

$$\sigma(1) = 12, \quad \sigma(2) = 31, \quad \sigma(3) = 34, \quad \sigma(4) = 13.$$

Set $u = 01101$ and $v = 001$; note that uv is a prefix of \mathbf{T} . Also define morphisms $g, h : \{1, 2, 3, 4\} \rightarrow \{0, 1\}^*$ by

$$g(1) = v\bar{u}, \quad g(2) = \bar{v}\bar{u}, \quad g(3) = \bar{v}u, \quad g(4) = vu$$

and

$$h(1) = uv, \quad h(2) = \bar{u}\bar{v}, \quad h(3) = \bar{u}\bar{v}, \quad h(4) = uv$$

We will make use of the following lemmas:

Lemma 1. $g(\sigma(a)) \in \mathcal{S}(g(a), h(a))$ for each $a \in \{1, 2, 3, 4\}$. In particular $ug(\sigma(1)) \in \mathcal{S}(ug(1), h(1))$.

Proof. For $a = 1$ we note that

$$g(\sigma(1)) = g(12) = v\bar{u}\bar{v}\bar{u} = 0011001011010010.$$

Factoring $0011001011010010 = 0 \cdot 011 \cdot 0 \cdot 010 \cdot 11 \cdot 01 \cdot 0010$ we obtain

$$g(\sigma(1)) \in \mathcal{S}(00110010, 01101001) = \mathcal{S}(v\bar{u}, uv) = \mathcal{S}(g(1), h(1)).$$

Similarly, for $a = 2$ we have

$$g(\sigma(2)) = g(31) = \bar{v}uv\bar{u} = 1100110100110010.$$

Factoring $1100110100110010 = 1 \cdot 100 \cdot 1 \cdot 1 \cdot 010 \cdot 0110 \cdot 010$ we obtain

$$g(\sigma(2)) \in \mathcal{S}(11010010, 10010110) = \mathcal{S}(\bar{v}\bar{u}, \bar{u}\bar{v}) = \mathcal{S}(g(2), h(2)).$$

Exchanging 0s and 1s in the previous two shuffles yields

$$g(\sigma(3)) = g(34) = \bar{v}uvu \in \mathcal{S}(\bar{v}u, \bar{u}\bar{v}) = \mathcal{S}(g(3), h(3))$$

and

$$g(\sigma(4)) = g(13) = v\bar{u}\bar{v}u \in \mathcal{S}(vu, uv) = \mathcal{S}(g(4), h(4)).$$

It is readily verified that

Lemma 2. $h(\sigma(a)) = \tau(h(a))$ for each $a \in \{1, 2, 3, 4\}$.

Let $w = w_0w_1w_2w_3 \dots$ with $w_i \in \{1, 2, 3, 4\}$ denote the fixed point of σ beginning in 1. As a consequence of the previous lemma we deduce that

Lemma 3. $\mathbf{T} = h(w)$.

Proof. In fact $\tau(h(w)) = h(\sigma(w)) = h(w)$ from which it follows that $h(w)$ is one of the two fixed points of τ . Since $h(w)$ begins in $h(1)$ which in turn begins in 0, it follows that $\mathbf{T} = h(w)$.

Lemma 4. $\mathbf{T} = ug(w)$.

Proof. It is readily verified that:

$$ug(1) = h(1)\bar{u}$$

$$\bar{u}g(2) = h(2)\bar{u}$$

$$\bar{u}g(3) = h(3)u$$

$$ug(4) = h(4)u.$$

Moreover, each occurrence of $g(1)$ and $g(4)$ in $ug(w)$ is preceded by u while each occurrence of $g(2)$ and $g(3)$ in $ug(w)$ is preceded by \bar{u} . It follows that $ug(w) = h(w)$ which by the preceding lemma equals \mathbf{T} .

Set

$$A_0 = ug(\sigma(w_0)) \quad \text{and} \quad A_i = g(\sigma(w_i)), \quad \text{for } i \geq 1$$

$$B_0 = ug(w_0) \quad \text{and} \quad B_i = g(w_i), \quad \text{for } i \geq 1$$

and

$$C_i = h(w_i) \quad \text{for } i \geq 0.$$

It follows from Lemma 3 and Lemma 4 that

$$\mathbf{T} = \prod_{i=0}^{\infty} A_i = \prod_{i=0}^{\infty} B_i = \prod_{i=0}^{\infty} C_i$$

and it follows from Lemma 1 that $A_i \in \mathcal{S}(B_i, C_i)$ for each $i \geq 0$. Hence $\mathbf{T} \in \mathcal{S}(\mathbf{T}, \mathbf{T})$ as required.

5 Self-Shuffling Sturmian Words

In this section we characterize self-shuffling Sturmian words. Sturmian words admit various types of characterizations of geometric and combinatorial nature, e. g., they can be defined via balance, complexity, morphisms, etc. (see [8]). In [11], Hedlund and Morse showed that each Sturmian word may be realized geometrically by an irrational rotation on the circle. More precisely, every Sturmian word x is obtained by coding the symbolic orbit of a point $\rho(x)$ on the circle (of circumference one) under a rotation by an irrational angle α where the circle is partitioned into two complementary intervals, one of length α (labeled 1) and the other of length $1 - \alpha$ (labeled 0) (see Figure 1). And conversely each such coding gives rise to a Sturmian word. The irrational α is called the *slope* and the point $\rho(x)$ is called the *intercept* of the Sturmian word x . A Sturmian word x of slope α with $\rho(x) = \alpha$ is called a *characteristic Sturmian word*. It is well known that every prefix u of a characteristic Sturmian word is *left special*, i.e., both $0u$ and $1u$ are factors of x [8]. Thus if x is a characteristic Sturmian word of slope α , then both $0x$ and $1x$ are Sturmian words of slope α and $\rho(0x) = \rho(1x) = 0$. The fact that ρ is not one-to-one stems from the ambiguity of the coding of the boundary points 0 and $1 - \alpha$.




Fig. 1. Geometric picture of a Sturmian word of slope α .

Theorem 4. *Let S , M and L be Sturmian words of the same slope α , $0 < \alpha < 1$, satisfying $S \leq_{\text{lex}} M \leq_{\text{lex}} L$. Then $M \in \mathcal{S}(S, L)$ if and only if the following conditions hold: If $\rho(M) = \rho(S)$ (respectively, $\rho(M) = \rho(L)$), then $\rho(L) \neq 0$ (respectively $\rho(S) \neq 0$).*

In particular (taking $S = M = L$), we obtain

Corollary 3. *A Sturmian word $x \in \{0, 1\}^{\mathbb{N}}$ is self-shuffling if and only if $\rho(x) \neq 0$, or equivalently, x is not of the form aC where $a \in \{0, 1\}$ and C is a characteristic Sturmian word.*

Our proof explicitly describes an algorithm for shuffling S and L so as to produce M . It is formulated in terms of the circle rotation description of Sturmian words. Geometrically speaking, points $\rho(S)$ and $\rho(L)$ will take turns following the trajectory of $\rho(M)$ so that the respective codings agree; as one follows the other waits its turn (remains neutral). The algorithm specifies this following rule depending on the relative positions of the trajectories of all three points and is broken down into several cases. The proof can be summarized by the directed graph in Figure 2 in which each state n corresponds to “case n ” in the proof.




Fig. 2. Graphical depiction of the proof of Theorem 4.

We denote by s , m , and ℓ the current tail of the words S , M , and L . They are initialized as

$$s := S, \ell := L, \text{ and } m := M.$$

While m is always a tail of M , the letters s and ℓ may be tails of S or L , depending on which is the current lexicographically largest⁶. Each directed edge corresponds to a precise set of instructions which specify which of s or ℓ is neutral, which of s or ℓ follows m and for how long, and in the end a possible relabeling of the variables s and ℓ . In each case the outcome leads to a new case in which there is a switch in the follower. In other words, if there is an edge from case i to case j in the graph, then either the instructions for case i and case j specify different followers (as is the case for cases 1.1 and 2.1) in which case the passage from i to j leaves the labeling of s and ℓ unchanged, or the instructions for case i and case j specify the same follower (as is the case for cases 1.2 and 1.1) in which case the passage from i to j exchanges the labeling of s and ℓ . The proof of Theorem 4 amounts to showing that for each state n in the graph, the specified instructions will take n to an adjacent state in the graph.

As an almost immediate application of Corollary 3 we recover the following result originally proved by Yasutomi in [14] and later reproved by Berth  , Ei, Ito and Rao in [3] and independently by Fagnot in [6]. We say an infinite word is *pure morphic* if it is a fixed point of some morphism different from the identity.

Theorem 5 (Yasutomi [14]). *Let $x \in \{0, 1\}^{\mathbb{N}}$ be a characteristic Sturmian word. If y is a pure morphic word in the orbit of x , then $y \in \{x, 0x, 1x, 01x, 10x\}$.*

Proof. We begin with some preliminary observations. Let $\Omega(x)$ denote the set of all left and right infinite words y such that $\mathcal{F}(x) = \mathcal{F}(y)$ where $\mathcal{F}(x)$ and $\mathcal{F}(y)$ denote the set of all factors of x and y respectively. If $y \in \Omega(x)$ is a right infinite word, and $0y, 1y \in \Omega(x)$, then $y = x$. This is because every prefix of y is a left special factor and hence also a prefix of the characteristic word x . Similarly if y is a left infinite word and $y0, y1 \in \Omega(x)$, then y is equal to the reversal of x . If τ is a morphism fixing some point $y \in \Omega(x)$, then $\tau(z) \in \Omega(x)$ for all $z \in \Omega(x)$.

Suppose to the contrary that $\tau \neq id$ is a morphism fixing a proper tail y of x . Then y is self-shuffling by Corollary 3. Put $x = uy$ with $u \in \{0, 1\}^+$. Using the characterization of Sturmian morphisms (see Theorem 2.3.7 & Lemma 2.3.13 in [8]) we deduce that τ must be primitive. Thus we can assume that $|\tau(a)| > 1$ for each $a \in \{0, 1\}$. If $\tau(0)$ and $\tau(1)$ end in distinct letters, then as both $0\tau(x), 1\tau(x) \in \Omega(x)$, it follows that $\tau(x) = x$. Since also $\tau(y) = y$ and $|\tau(u)| > |u|$, it follows that y is a proper tail of itself, a contradiction since x is aperiodic. Thus $\tau(0)$ and $\tau(1)$ must end in the same letter. Whence by Corollary 1 it follows that every backward extension of y is self-shuffling, which is again a contradiction since $0x$ and $1x$ are not self-shuffling.

Next suppose $\tau \neq id$ is a morphism fixing a point $y = uabx \in \Omega(x)$ where $u \in \{0, 1\}^+$ and $\{a, b\} = \{0, 1\}$. Again we can suppose τ is primitive and $|\tau(0)| > 1$ and $|\tau(1)| > 1$. If $\tau(0)$ and $\tau(1)$ begin in distinct letters, then $\tau(\tilde{x})0, \tau(\tilde{x})1 \in \Omega(x)$ where \tilde{x} denotes the reverse of x . Thus $\tau(\tilde{x}) = \tilde{x}$. Thus for each prefix v of abx we have $\tau(\tilde{x}v) = \tilde{x}\tau(v)$ whence $\tau(v)$ is also a prefix of abx . Hence $\tau(abx) = abx$.

⁶ The choice of the letter s, m , and ℓ is intended to refer to *small*, *medium*, and *large* respectively.

As before this implies that abx is a proper tail of itself which is a contradiction. Thus $\tau(0)$ and $\tau(1)$ begin in the same letter. Whence by Corollary 1 it follows that every tail of y is self-shuffling, which is again a contradiction since $0x$ and $1x$ are not self-shuffling.

Remark 3. In the case of the Fibonacci infinite word x , each of $\{x, 0x, 1x, 01x, 10x\}$ is pure morphic. For a general characteristic word x , since every point in the orbit of x except for $0x$ and $1x$ is self-shuffling, it follows that if τ is a morphism fixing x (respectively $01x$ or $10x$), then $\tau(0)$ and $\tau(1)$ must end (respectively begin) in distinct letters.

References

1. J.-P. Allouche and J. Shallit, *The ubiquitous Prouhet-Thue-Morse sequence*, in: *Sequences and their applications*, Proceedings of SETA'98, C. Ding, T. Helleseth and H. Niederreiter (Eds.) (1999), Springer Verlag, 1–16.
2. J.-P. Allouche and J. Shallit, *Automatic sequences : Theory, applications, generalizations*, Cambridge University Press, 2003.
3. V. Berthé, H. Ei, S. Ito, H. Rao, On substitution invariant Sturmian words: an application of Rauzy fractals, *Theor. Inform. Appl.*, **41** (2007) 329–349.
4. J. Cassaigne, J. Karhumäki, Toeplitz Words, Generalized Periodicity and Periodically Iterated Morphisms, *European J. Combin.*, **18** (1997) 497–510.
5. D. Henshall, N. Rampersad, J. Shallit, Shuffling and unshuffling, *Bull. EATCS*, **107** (2012) 131–142.
6. I. Fagnot, A little more about morphic Sturmian words, *Theor. Inform. Appl.*, **40** (2006) 511–518.
7. M. Lothaire, *Combinatorics On Words*, vol. 17 of *Encyclopedia of Mathematics and its Applications*, Addison-Wesley, Reading, Massachusetts, 1983.
8. M. Lothaire, *Algebraic Combinatorics On Words, Chapter 2: Sturmian words*, by J. Berstel, P. Séébold, vol. 90 of *Encyclopedia of Mathematics and its Applications*, Cambridge University Press, U.K., 2002.
9. M. Morse, *Recurrent geodesics on a surface of negative curvature*, Trans. Amer. Math. Soc. **22** (1921), p. 84–100.
10. M. Morse and G. Hedlund, *Symbolic dynamics*, Amer. J. Math. **60** (1938), p. 815–866.
11. M. Morse, G.A. Hedlund, Symbolic dynamics II: Sturmian sequences, *Amer. J. Math.* **62** (1940), p. 1–42.
12. A. Thue, *Über unendliche Zeichenreihen*, Norske Vid. Selsk. Skr. I Math-Nat. Kl. **7** (1906), 1–22.
13. A. Thue, *Die Lösung eines Spezialfalles eines generellen logischen Problems*, Norske Vid. Selsk. Skr. I Math-Nat. Kl. Chris. **8** (1910).
14. S.-I. Yasutomi, On sturmian sequences which are invariant under some substitutions, in *Number theory and its applications*, Proceedings of the conference held at the RIMS, Kyoto, Japan, November 1014, 1997, edited by Kanemitsu, Shigeru et al. Kluwer Acad. Publ. Dordrecht (1999) 347–373.

6 Appendix

6.1 Proof of Theorem 1

We state and prove the theorem in greater generality for k -self-shuffle.

Definition 3. For $x \in \mathbb{A}^{\mathbb{N}}$ and integer $k \geq 2$, define a directed graph $G_x^k = (V_x^k, E_x^k)$ with the vertex set

$$V_x^k = \{(i_1, \dots, i_k) \in \mathbb{N}^k; \\ \sum_{j=1}^k \Psi(\text{pref}_{i_j} x) = \Psi(\text{pref}_{i_1+\dots+i_k} x)\},$$

and the edge set

$$E_x^k = \{((i_1, \dots, i_k), (i'_1, \dots, i'_k)) \in V_x^k \times V_x^k; \\ i_j \leq i'_j \text{ for } j = 1, \dots, k \text{ and } \sum_j i_j + 1 = \sum_j i'_j\}.$$

We say that G_x^k connects $\mathbf{0}$ to ∞ if there exists an infinite path $v^0 v^1 v^2 \dots$ in G_x^k such that $v^0 = (0, \dots, 0)$ and $v_j^n \rightarrow \infty$ as $n \rightarrow \infty$ for any $j = 1, \dots, k$, where $v^n = (v_1^n, \dots, v_k^n) \in V_x^k$.

Theorem 6. An infinite sequence $x \in \mathbb{A}^{\mathbb{N}}$ is k -self-shuffling if and only if the graph G_x^k connects $\mathbf{0}$ to ∞ .

Proof. If $x = x_0 x_1 \dots$ with the $x_i \in A$ is k -self-shuffling, then there exist infinite subsets $N^j \subset \mathbb{N}$ for $j = 1, \dots, k$, such that $x[N^j] = x$. Let $N^j = \{N_0^j < N_1^j < N_2^j < \dots\}$ ($j = 1, \dots, k$). For any $n \in \mathbb{N}$, let $t(j, n)$ be $m \in \mathbb{N}$ such that $N_{m-1}^j < n \leq N_m^j$ for any $j = 1, \dots, k$. Then, $\text{pref}_n x$ consists of $\text{pref}_{t(j,n)} x[N^j]$ with $j = 1, \dots, k$ in Abelian sense. Since

$$\text{pref}_{t(j,n)} x[N^j] = \text{pref}_{t(j,n)} x \quad (j = 1, \dots, k),$$

we have

$$\sum_{j=1}^k \Psi(\text{pref}_{i_j} x) = \Psi(\text{pref}_n x)$$

with $i_j = t(j, n)$ and $i_1 + \dots + i_k = n$. Therefore, $(i_1, \dots, i_k) \in V_x^k$. Let $n \in N^l$ for some $l = 1, \dots, k$ and $x_n = b \in \mathbb{A}$, then we have

$$i'_j := t(j, n+1) = \begin{cases} t(j, n) + 1 & \text{for } j = l \\ t(j, n) & \text{for } j \neq l \end{cases},$$

and

$$\Psi(\text{pref}_{i'_j} x[N^j]) = \begin{cases} \Psi(\text{pref}_{i_j} x[N^j]) + e_b & \text{for } j = l \\ \Psi(\text{pref}_{i_j} x[N^j]) & \text{for } j \neq l \end{cases},$$

where e_b is a unit vector having 1 in a coordinate corresponding to the letter b . Moreover, since $\Psi(\text{pref}_{n+1} x) = \Psi(\text{pref}_n x) + e_b$, we have

$$\sum_{j=1}^k \Psi(\text{pref}_{i'_j} x) = \Psi(\text{pref}_{n+1} x).$$

Thus, $((i_1, \dots, i_k), (i'_1, \dots, i'_k)) \in E_x^k$. Since this holds for any $n \in \mathbb{N}$ and each N^i is an infinite set, G_x^k connects $\mathbf{0}$ to ∞ .

Conversely, assume that G_x^k connects $\mathbf{0}$ to ∞ . Let $v^0 v^1 v^2 \dots$ be an infinite path in G_x^k such that $v^0 = (0, \dots, 0)$ and $v_j^n \mapsto \infty$ as $n \mapsto \infty$ for any $j = 1, \dots, k$, where $v^n = (v_1^n, \dots, v_k^n) \in V_x^k$. Define

$$N^j = \{n \in \mathbb{N}; v_j^{n+1} > v_j^n\} \text{ for } j = 1, \dots, k.$$

Then the sets N^j give a partition of \mathbb{N} . Let $N^j = \{N_0^j < N_1^j < N_2^j < \dots\}$ for $j = 1, \dots, k$. For any $l = 1, \dots, k$ and $m \in \mathbb{N}$, let $n = N_m^l$, $i_j = t(j, n)$ and $i'_j = t(j, n+1)$ for $j = 1, \dots, k$. Then, $v^n = (i_1, \dots, i_k)$ and $v^{n+1} = (i'_1, \dots, i'_k)$. Moreover, $i_l = m$ and $i'_l = m+1$. Since $(v^n, v^{n+1}) \in E_x^k$, we have

$$\begin{aligned} \Psi(\text{pref}_{m+1}x) - \Psi(\text{pref}_m x) &= \sum_{j=1}^k (\Psi(\text{pref}_{i'_j} x) - \Psi(\text{pref}_{i_j} x)) \\ &= \Psi(\text{pref}_{n+1}x) - \Psi(\text{pref}_n x). \end{aligned}$$

Hence, $x_m = x_n = x[N^l]_m$. Thus, the sets N^j satisfy Definition 1, so we have a k -self-shuffle.

6.2 Proof of Theorem 2

Here we sketch the proof of Theorem 2.

Proof. We will make use of the following four lemmas whose proofs are omitted:

Lemma 5. *Let $u = u_1 u_2 \dots u_n$ be a factor of x with each u_i minimal in x . Then u is minimal in x .*

A word $u \in \mathbb{A}^+$ is called *bordered*, if there exists an integer k , $0 < k < |u|$ such that $\text{pref}_k u = \text{suff}_k u$. Otherwise u is called *unbordered*.

Lemma 6. *Let u be a minimal factor of x and let v be the longest unbordered prefix of u . Then v is a period of u , i.e., u is a prefix of v^n for some positive integer n .*

Lemma 7. *Let u and v be factors of $x \in \mathbb{A}^{\mathbb{N}}$ with u minimal. Then either $uv <_{\text{lex}} v$ or else v is minimal.*

Let $\mathcal{C}(x)$ denote the set of all factors v of x with the property that no suffix of v (including v itself) is minimal in x .

Lemma 8. *Let u and v be factors of x with u minimal in x and $v \in \mathcal{C}(x)$. Let $s \in \mathcal{S}(u, v)$. Then either $s = vu$ or $s <_{\text{lex}} v$.*

We first consider the case where $z = y$, i.e, when $w \in \mathcal{S}(y, y)$. Set

$$w = \prod_{i=0}^{\infty} U_i V_i$$

where

$$y = \prod_{i=0}^{\infty} U_i = \prod_{i=0}^{\infty} V_i$$

with each U_i and V_i non-empty. Since y is aperiodic and Lyndon, it follows from Lemma 6 that y contains arbitrarily long unbordered prefixes. Let v be an unbordered prefix of y with $|v| > |U_0|$. Writing $v = U_0 v'$, since v is unbordered, $v' \in \mathcal{C}(x)$. Let s be such that $U_0 s$ is a prefix of w and $|s| = |v'|$. Then by Lemma 8 we deduce that $s <_{\text{lex}} v'$ whence $U_0 s <_{\text{lex}} v$ and hence $w <_{\text{lex}} y$.

Next suppose $z \neq y$. Let

$$w = \prod_{i=0}^{\infty} U_i V_i$$

where

$$y = \prod_{i=0}^{\infty} U_i \quad \text{and} \quad z = \prod_{i=0}^{\infty} V_i$$

with each U_i and V_i non-empty except for possibly U_0 . Suppose first that U_0 is non-empty, that is to say y dishes out the initial segment of w . Let r be the shortest non-minimal prefix of z . Then by Lemma 5 $r \in \mathcal{C}(x)$. Let u be a prefix of y longer than $|U_0|$. Let t denote the prefix of w of length $|r|$; then by Lemma 8, $t <_{\text{lex}} r$ and hence $w <_{\text{lex}} z$. Finally suppose that U_0 is empty, so that z dishes out the initial segment of w . Let z' be a tail of z so that $z = V_0 z'$. Writing $w = V_0 w'$, it follows that w' is a shuffle of z' and y in which y dishes out the initial segment of z' . If $z' \neq y$ then we are done by the preceding case, i.e, $w' <_{\text{lex}} z'$ whence $w <_{\text{lex}} z$. If $z' = y$, then as we saw in the beginning of the proof, we again have $w' <_{\text{lex}} y = z'$ whence $w <_{\text{lex}} z$.

6.3 Proof of Theorem 4

We begin by showing that the conditions stated in Theorem 4 are in fact necessary for $M \in \mathcal{S}(S, L)$. To see this, suppose $\rho(M) = \rho(S)$ and $\rho(L) = 0$ (the other symmetric condition works analogously). This implies that $L \in \{0x, 1x\}$ where x is the characteristic Sturmian word of slope α . If $L = 0x$, then as $0x$ is minimal in the Sturmian subshift of slope α , it follows that $S = M = L$. Whence by Corollary 2, $M \notin \mathcal{S}(M, M) = \mathcal{S}(S, L)$. If $L = 1x$, we consider the lexicographic order induced by $0 > 1$. Then $L \leq_{\text{lex}} M \leq_{\text{lex}} S$ and moreover L is minimal. Since $\rho(M) = \rho(S)$ we have that either case i) $M = S$ or case ii) $S = 0x$ and $M = 1x$ or case iii) there exists $u \in \{0, 1\}^*$ such that $S = u01x$ and $M = u10x$ where in each case x denotes the characteristic Sturmian word of slope α . In case i), using Theorem 2 we deduce that each element of $\mathcal{S}(S, L)$ is

lexicographically smaller than S and hence since $M = S$ we have $M \notin \mathcal{S}(S, L)$. In case ii), if $M \in \mathcal{S}(S, L)$, then $x \in \mathcal{S}(x, 0x)$ which contradicts Theorem 2. Finally, in case iii), suppose to the contrary that $M \in \mathcal{S}(S, L)$. Then since $u0$ is not a prefix of M , it follows that there exists a non-empty prefix v of L and a prefix w of M such that $|w| = |u0| + |v|$ and $M' \in \mathcal{S}(L', 1x)$ where M' and L' are defined by $M = wM'$ and $L = vL'$. But this implies that $M' = L'$, whence $L' \in \mathcal{S}(L', 1x)$ which again contradicts Theorem 2.

We next show that the conditions stated in Theorem 4 are sufficient. Without loss of generality we can assume that $0 < \alpha < 1/2$.

Our proof explicitly describes an algorithm for shuffling S and L so as to produce M . It is formulated in terms of the circle rotation description of Sturmian words. Geometrically speaking, points $\rho(S)$ and $\rho(L)$ will take turns following the trajectory of $\rho(M)$ so that the respective codings agree; as one follows the other waits its turn (remains neutral). The algorithm specifies this following rule depending on the relative positions of the trajectories of all three points and is broken down into several cases. The proof can be summarized by the directed graph in Figure 2 in which each state n corresponds to “case n ” in the proof.

The states : We denote by s , m , and ℓ the current tail of the words S , M , and L . They are initialized as

$$s := S, \ell := L, \text{ and } m := M.$$

While m is always a tail of M , the letters s and ℓ may be tails of S or L , depending on which is the current lexicographically largest. Each state, or case in the proof, is described by a figure depicting the relative positions of $\rho(s)$, $\rho(m)$ and $\rho(\ell)$, which for the sake of simplicity, are actually labeled s , m and ℓ respectively. If $x \in \{s, m, \ell\}$ is depicted inside the interval $(0, 1 - \alpha)$ (resp. $(1 - \alpha, 1)$), then this implies that the first letter of the coding of x is 0 (respectively 1). Moreover the endpoints of the partition interval are regarded as both open and closed. For example, even if s and m are both depicted in the interval $(0, 1 - \alpha)$, it could be that $\rho(s) = 0$ and $\rho(m) = 1 - \alpha$. In the same way, even if s and m are depicted in distinct intervals of the circle partition, it could be that $\rho(s) = \rho(m)$. In addition to their relative positions on the circle, each state lists a set of relations which the variables s , m and ℓ must satisfy. These conditions are written to the right of the circle picture and are described in terms of the following predicates:

$$C(s, m, \ell) \equiv [\rho(m) = \rho(s) \text{ and } \rho(\ell) = 0] \text{ or } [\rho(m) = \rho(\ell) \text{ and } \rho(s) = 0];$$

$$P_1(s, m, \ell) \equiv \rho(s) = \alpha \text{ and } \rho(m) = 0 \text{ and } \rho(\ell) = 1 - \alpha;$$

$$P_2(s, m, \ell) \equiv [(\rho(\ell) - \rho(m)) \bmod 1 = \alpha \text{ and } \rho(s) = 1 - \alpha] \text{ or } \rho(m) = 0.$$

All states, except those labeled 4 and 5, can be taken to be initial states.


The edges : Each directed edge corresponds to a precise set of instructions which specify which of s or ℓ is neutral, which of s or ℓ follows m and for how long, and in the end a possible relabeling of the variables s and ℓ . In each case the outcome leads to a new case in which there is a switch in the follower. In

other words, if there is an edge from case i to case j in the graph, then either the instructions for case i and case j specify different followers (as is the case for cases 1.1 and 2.1) in which case the passage from i to j leaves the labeling of s and ℓ unchanged, or the instructions for case i and case j specify the same follower (as is the case for cases 1.2 and 1.1) in which case the passage from i to j exchanges the labeling of s and ℓ .

The proof of Theorem 4 amounts to showing that for each state n in the graph, the specified instructions will take n to an adjacent state in the graph.


What follows is a complete listing of all ten cases with their respective set of instructions.

Case 1.1:




Instruction: s is neutral and ℓ follows m until they lie in different elements of the circle partition. No relabeling of s and ℓ .

Case 1.2:




Instruction: s is neutral and ℓ follows m until $0 < \rho(m) = \rho(\ell) < \rho(s)$. We note that this is always possible since $\rho(s) \neq 0$ and the set $\{(\rho(m) + n\alpha) \bmod 1 : n \in \mathbb{N}\}$ is dense in the unit circle. Exchange the labels $s \leftrightarrow \ell$.

Case 2.1:



Instruction: ℓ is neutral and s follows m until they lie in different elements of the circle partition. No relabeling of s and ℓ . Three cases are possible according to the relative position of m and ℓ in the partition $(1 - \alpha, 1)$.


Case 2.2:



$$s = m; \neg C(s, m, \ell)$$

Instruction: ℓ is neutral and s follows m until $\rho(m) = \rho(s) > \rho(\ell)$. This is possible because $\rho(\ell) \neq 0$. Exchange the labels $s \leftrightarrow \ell$.


Case 3.1:



$$m <_{\text{lex}} \ell; \neg C(s, m, \ell)$$

Instruction: s is neutral and ℓ follows m until they lie in different elements of the circle partition. No relabeling of s and ℓ . Three cases are possible according to the relative position of m and s in the partition $(0, 1 - \alpha)$.


Case 3.2:



$$m = \ell; \neg C(s, m, \ell)$$

Instruction: s is neutral and ℓ follows m until $0 < \rho(m) = \rho(\ell) < \rho(s)$. This is possible because $\rho(s) \neq 0$. Exchange the labels $s \leftrightarrow \ell$.


Case 4:



$$m >_{\text{lex}} \ell; \rho(s) \geq \alpha; \neg P_1(s, m, \ell)$$


Instruction: s is neutral and ℓ follows m for just one rotation by α . Exchange the labels $s \leftrightarrow \ell$. Because $\rho(s) \geq \alpha$, we have either $m <_{\text{lex}} s$ or $m = s$.

Case 5:




Instruction: ℓ is neutral and s follows m for just one rotation by α . Exchange the labels $s \leftrightarrow \ell$. Because $(\rho(\ell) - \rho(m)) \bmod 1 \leq \alpha$, we have either $m >_{\text{lex}} \ell$ or $m = \ell$.

Case 6.1:



Instruction: ℓ is neutral and s follows m until they lie in different elements of the circle partition. No relabeling of s and ℓ .

Case 6.2:



Instruction: ℓ is neutral and s follows m until: $\rho(m) = \rho(s) > \rho(\ell)$. This is possible because $\rho(\ell) \neq 0$. Exchange the labels $s \leftrightarrow \ell$.

Here we verify four of the ten cases in the proof of Theorem 4. The verifications of all cases are similar to one another.

Verification of Case 1.1: To see that Case 1.1 leads to Case 2.1, let m' and ℓ' denote the positions of m and ℓ respectively, the first time they lie in different elements of the circle partition. Then clearly $0 \leq \rho(s) < \rho(m') \leq 1 - \alpha \leq \rho(\ell')$ as required. It remains to show that after the rotation $\neg C(s, m', \ell')$ holds. Suppose to the contrary that $C(s, m', \ell')$ holds. Because $\rho(m') > \rho(s)$, we must have $\rho(m') = \rho(\ell')$ and $\rho(s) = 0$. But this implies $\rho(m) = \rho(\ell)$ and $\rho(s) = 0$, which is impossible since we had assumed $\neg C(s, m, \ell)$.

Verification of Case 1.2: To see that Case 1.2 leads to Case 1.1, let m' and ℓ' denote the positions of m and ℓ respectively, the first time $0 < \rho(m') =$

$\rho(\ell') < \rho(s)$. Then clearly after exchanging the labels $s \leftrightarrow \ell$ the points s, m , and ℓ are situated as specified in Case 1.1. It remains to show that $\neg C(\ell', m', s)$. Suppose to the contrary that $C(\ell', m', s)$ holds. Because $m' = \ell'$, we must have $\rho(m') = \rho(\ell')$ and $\rho(s) = 0$. But this implies $\rho(m) = \rho(\ell)$ and $\rho(s) = 0$, which is impossible since we had assumed $\neg C(s, m, \ell)$.

Verification of Case 2.1: Let s' and m' denote the positions of s and m respectively, the first time they lie in different elements of the circle partition. Note that since we have assumed $\alpha < 1/2$, it follows that $\rho(s') \geq \alpha$ for otherwise m and s would have already differed earlier. Three cases are possible: $m' <_{\text{lex}} \ell$, $m' = \ell$ or $m' >_{\text{lex}} \ell$. We show that this leads to cases 3.1, 3.2 and 4 respectively. Assume first that $m' \leq_{\text{lex}} \ell$. To show that this leads to Case 3.1 or Case 3.2, we must verify that $\neg C(s', m', \ell)$. However we have $\alpha \leq \rho(s') \leq 1 - \alpha$, and hence $\rho(s') \neq 0$. If $\rho(m') = \rho(s')$, then $\rho(m) = \rho(s)$, and hence $\rho(\ell) \neq 0$ since we had assumed $\neg C(s, m, \ell)$. Next we suppose that $m' > \ell$. To show that this results in Case 4, we must show that

$$\neg P_1(s', m', \ell).$$

Assume to the contrary that $P_1(s', m', \ell)$, that is, that $\rho(s') = \alpha$, $\rho(m') = 0$ and $\rho(\ell) = 1 - \alpha$. This implies $m = 0m'$ and $s = 0s'$, and hence $\rho(m) = 1 - \alpha = \rho(\ell)$ and $\rho(s) = 0$, which is impossible since we had assumed $\neg C(s, m, \ell)$.

Verification of Case 4: Let ℓ' and m' denote the positions of ℓ and m after rotation by α . Because $\rho(s) \geq \alpha$, we have either $m' <_{\text{lex}} s$ or $m' = s$. We will show that this leads to cases 1.1 and 1.2 respectively. In view of the label exchange $s \leftrightarrow \ell$, the relative positions of the three points is as required. It remains to check in both cases that $\neg C(\ell', m', s)$ holds. Since $\rho(s) \geq \alpha$, it follows that $\rho(s) \neq 0$. If $\rho(m') = \rho(s)$, then we actually have $\rho(m') = \rho(s) = \alpha$. This implies $\rho(m) = 0$. Because we had assumed $\neg P_1(s, m, \ell)$, we obtain $\rho(\ell) \neq 1 - \alpha$, and hence $\rho(\ell') \neq 0$ as required.